On the capacity of a new model of associative memory based on neural cliques

Judith Heusel, Matthias Löwe, Franck Vermet

arXiv:1411.1224v1 [math.PR] 5 Nov 2014

Abstract— **Based on recent work by Gripon and Berrou, we introduce a new model of an associative memory. We show that this model has an efficiency bounded away from 0 and is therefore significantly more effective than the well known Hopfield model. We prove that the synchronous and asynchronous retrieval dynamics converge and give upper and lower bounds on the memory capacity of the model.**

I. Introduction

In [6] Gripon and Berrou introduced a new model of an associative memory. This model seems to be more effective than standard models of associative memories, in particular the Hopfield model, while at the same time it respects the basic principles of associative memories, e.g. locality. In this model, which we will call GB-model for short, there are N neurons grouped into c groups of l neurons. Typically, one c is bounded from below by a constant and from above by $\log l$, where the latter bound seems to be the most interesting case. One tries to store M messages $m^1, \dots, m^M$ in this network. These messages are sparse in the sense that each message m^μ has c active neurons, only, i.e. $m^\mu = (m_1^\mu, \dots, m_c^\mu)$, and for each $\mu = 1, \dots M$ and each $i = 1, \dots c$, m_i^μ denotes the (only) active neuron of the message m_μ in the i'th block. With such a message m^μ one associates the edges of the complete graph K^μ spanned by the vertices $m_1^\mu, \dots, m_c^\mu$. A message $m^0 = (m_1^0, \dots, m_c^0)$ is considered to be stored in the model, if all edges of the complete graph spanned by $(m_1^0, \dots, m_c^0)$ are present in the set of edges

$$\mathcal{E} := \{e : e \text{ is an edge of one of the } m^\mu\}.$$

Gripon and Berrou now analyze in a number of papers (see e.g. [6],[1], [2] among others) the performance of this network for a random input $m^1, \dots, m^M$, e.g. they ask for the probability that a random pattern m^0 is stored or that a corrupted pattern can be retrieved. However, all their analysis is either based on numerical simulations or on the assumption that the events that a given edge occurs in $\mathcal{E}$ are independent. This is, of course, not true, even though the dependence is very weak. On the other hand, the precise form of the network makes it difficult to analyze it rigorously. In the present paper we therefore strive for the rigorous analysis of yet another associative memory model, which is, however, closely related to the GB model.

To motivate it, observe that the GB model can be mathematically described in the following way: We start with an alphabet $\mathcal{A} = \{1, \dots, l\}$. A message m^μ is then a string $m^\mu = (m_1^\mu, \dots, m_c^\mu) \in \mathcal{A}^c$. With a message m^μ we associate a (column) vector $\psi(m^\mu) \in (\{0,1\}^l)^c$ obtained by replacing the m_i^μ with the unit vector $e_{m_i^\mu}$. In a slight abuse of notation we will also use the notation $\mathcal{A}^c$ for the set $\{e_1, \dots, e_l\}^c$. Now build the the 0-1-matrix W given by

$$\tilde{W} = \max_{m \in \mathcal{M}} \psi(m)\psi(m)^t$$

where $\mathcal{M} = \{m^1, \dots, m^M\}$ and $\psi(m)^t$ is the transpose of $\psi(m)$. Thus

Fachbereich Mathematik und Informatik, University of Münster, Einsteinstraße 62, 48149 Münster, email: jheus01@uni-muenster.de

Institut für Mathematische Statistik, Fachbereich Mathematik und Informatik, Universität Münster, Einsteinstr. 62, 48149 Münster, Germany, email: maloewe@math.uni-muenster.de

Departement de Mathematiques, Universite de Bretagne Occidentale, 6, avenue Victor Le Gorgeu B.P. 809 F- 29285 Brest Cedex, France, email: Franck.Vermet@univ-brest.fr

- for $a \neq a'$ we have $\tilde{W}_{(a,k),(a',k')} = 1$ if and only if there is an edge between (a,k) and (a',k').
- for $k \neq k'$ we have $\tilde{W}_{(a,k),(a,k')} = 0$
- and $\tilde{W}_{(a,k),(a,k)} = 1$ if and only if there exists μ such that the k'th neuron in block a is 1 (this is equivalent to adding a self-loop to the graph for each vertex (a,k) such that there exists μ with $m_a^\mu = k$.

With this matrix one can associate a dynamics D on $(\{0,1\}^l)^c$: for $v \in (\{0,1\}^l)^c$,

$$D(v)_{(a,k)} = \bigwedge_{b=1}^{c} \bigvee_{r=1}^{l} \tilde{W}_{(a,k),(b,r)} v_{(b,r)},$$

which can also be written as

$$D(v)_{(a,k)} = \mathbf{1}_{\{\sum_{b=1}^{c} \mathbf{1}_{\{\sum_{r=1}^{l} \tilde{W}_{(a,k),(b,r)} v_{(b,r)} \geq 1\}} \geq c\}}.$$

Obviously, for all learned message $m \in \mathcal{M}$, we have

$$D(\psi(m)) = \psi(m).$$

However, a more detailed analysis of the associative abilities of the network turn out to be difficult, basically due to the double indicator structure of D and the function max in $\tilde{W}$.

We therefore propose the following variant of the above model. Consider the Bernoulli random variables $\zeta^\mu_{(a,i)}$, that denote if neuron i of cluster a is part of message μ:

$$\zeta^\mu_{(a,i)} = \begin{cases} 1 & \text{if } m_a^\mu = e_i, \\ 0 & \text{otherwise.} \end{cases}$$

Under the above assumptions on the patterns, i.e. their independence and equidistribution, the random variables $(\zeta^\mu_{(a,i)})^{1 \leq \mu \leq M}_{1 \leq a \leq c, 1 \leq j \leq l}$ are Bernoulli variables, each with parameter $\frac{1}{l}$, and $\zeta^\mu_{(a,i)}$ is independent of $\zeta^\nu_{(b,j)}$ if $a \neq b$ or $\mu \neq \nu$.
We define the matrix $W \in \mathbb{N}^{cl \times cl}$ by

$$W_{(a,i),(b,j)} = \sum_{\mu=1}^{M} \zeta^\mu_{(a,i)} \zeta^\mu_{(b,j)},$$

for $a, b \in \{1, \ldots, c\}, a \neq b$ and $i, j \in \{1, \ldots, l\}$.

We set $W_{(a,i),(a,j)} = 0$ for all $i, j \in \{1, \ldots, l\}$ and $a \in \{1, \ldots, c\}$. Hence, $W_{(a,i),(a,i)} = 0$, while $\tilde{W}_{(a,i),(a,i)} \in \{0,1\}$. However, this choice is necessary due to the structure of W as a sum or random variables.

Given an input vector $v = (v_{(b,j)})_{1 \leq b \leq c, 1 \leq j \leq l} \in \{0,1\}^{cl}$, we define the dynamics

$$\varphi_{(a,i)}(v) = \mathbb{1}_{\{\sum_{b=1}^{c} \sum_{j=1}^{l} W_{(a,i),(b,j)} v_{(b,j)} \geq \kappa c\}} \tag{1}$$

for some $\kappa > 0$. It should be obvious that this variant is closely related to the GB model. On the other the structure of W and the dynamics that include sums of random variables rather maxima and minima, makes it more accessible to probabilistic tools. Indeed from this point of view the model is reminiscent to the Hopfield model, whose the storage capacity has been analyzed in [15], [9], [16], [10], [17], [11], [13],[14], and many other papers.

In the present note we want to justify our variant of the GB model by showing that it has a non-vanishing efficiency in the range of parameters we want to study. This will be done in Section 2. In Section 3 we introduce the sequential and parallel retrieval of the model and show that it converges, possibly to a cycle of length 2. Section 4 contains our bounds on the memory of storage capacity of our version of the GB model. Sections 5 and 6 are devoted to the proofs of these results.

II. The efficiency of the model

In this section we want to justify that our version of the GB model is potentially better than other models of neural networks, in particular the Hopfield model, that may be the best studied model of an associative memory. We will do this with a performance measure that is known as the *informational efficiency* η. η compares the amount of information in the memorized messages to the logarithm of the number of bits needed to describe all the bonds in a state of the network at the storage capacity.

There are good reasons for this definition (and not comparing the amount of information in the memorized messages to the information in the network, e.g.). First, for any reasonable neural network the state of the network is a function of the information to be stored. In particular, if we compared the information in

the messages to the information in the network and we were in the retrieval regime η would be 1 for any network. Second, the point in the entire theory of associative memories is that we do not know the stored information. Hence the best way to describe our network is to describe it bond by bond, as if they were independent (for a related discussion see also [8]).

For our network, in view of Section 4, we will always choose $M = \alpha l^2$ and $c = \log l$.

Recall that in information theory a source is a probability distribution $\pi = (\pi_1, \ldots, \pi_A)$ on a (finite) alphabet $\{1, \ldots, A\}$ (for background material the reader is referred to classical textbook of Ash [3]). Its information content is measured by the (binary) entropy

$$H(\pi) = -\sum_{i=1}^{A} \pi_A \log_2 \pi_A.$$

$H(\cdot)$ is maximized by the uniform distribution $u(i) = 1/A$ for all i, in which case we obtain $H(u) = \log_2 A$. Assuming that we sample the M words $\mathcal{M}$ with the uniform distribution on all possible pairwise different $\binom{l^c}{M}$ we obtain that

$$H(\mathcal{M}) = \log_2 \left(\binom{l^c}{M} \right) \simeq Mc \log_2(l),$$

for l large, $M = \alpha l^2$, and $c = \log(l)$.

On the other hand, the network is completely described by the symmetric matrix W. There are $\binom{c}{2} l^2$ many matrix entries. Moreover for $l \to \infty$ a single matrix entry is asymptotically Poisson-distributed with parameter α.

This implies that the entropy of the network satisfies

$$\begin{aligned} H(W) &= \binom{c}{2} l^2 H(W_{(1,1),(2,1)}) \\ &\simeq \binom{c}{2} l^2 \; h(\mathcal{P}\text{ois}(\alpha)), \end{aligned}$$

where

$$\begin{aligned} h(\mathcal{P}\text{ois}(\alpha)) &= \frac{1}{\log(2)} (\alpha(1 - \log(\alpha)) \\ &\quad + e^{-\alpha} \sum_{k=0}^{+\infty} \alpha^k \frac{\log(k!)}{k!}) \end{aligned}$$

is the entropy of the Poisson law with parameter α.

Therefore we obtain

$$\eta = \frac{H(\mathcal{M})}{H(W)} \simeq \frac{2\alpha}{\log(2) \; h(\mathcal{P}\text{ois}(\alpha))} > 0.$$

Note that in the Hopfield model with $M = \text{const.N}/\log \text{N}$ patterns the efficiency tends to 0 when N is large, while the efficiency for our model is bounded away from zero. Indeed, for the Hopfield model, $H(\mathcal{M})$ is of order $N^2/\log(N)$, while $H(W)$ is larger than N^2, since the matrix W has $N(N-1)/2$ different entries and each entry W_{ij} follows the Binomial law $\mathcal{B}in(M, 1/2)$.

Moreover also notice, that η is increasing with α, and one can compute numerically that $\eta > 1$ if $\alpha \geq 0.423$. Interestingly, in Theorem IV.5 we obtain similar critical value for α: Both results show that if α is too large, the system looses information in the sense that $H(\mathcal{M}) > H(W)$ and not all the information contained in the words can be stored in the system.

III. The retrieval dynamics

Already in (1) we introduced the dynamics φ of our version of the GB model.

As a matter of fact, one can associate either a sequential or a parallel dynamics with φ. In this section will prove that for all initial condition, both dynamics converge.

First let

$$S = \varphi_{(c,l)} \circ \varphi_{(c,l-1)} \circ \ldots \circ \varphi_{(1,1)}.$$

For an arbitrary initial state $v(0)$ we define the sequential dynamics $v(t+1) = S(v(t))$.

Then the following result holds:

Proposition III.1: Let $H = H_S$ be the following function:

$$H_S(v) = -\frac{1}{2} \sum_{(a,i)} \sum_{(b,j)} W_{(a,i),(b,j)} v_{(a,i)} v_{(b,j)} + \kappa c \sum_{(a,i)} v_{(a,i)}.$$

The reason, why it carries an index S is, that with H as Hamiltonian (or energy function), the dynamics S is a Hamiltonian zero temperature (or gradient descent) dynamics.

Then for all M, c, l, and $v(0)$,

$$H_S(v(t+1)) \leq H_S(v(t)) \qquad \text{for all } t,$$

and the sequential dynamics converges to a fixed point.

Proof: To simplify notation, we denote by capital letters the double indices, for example: $I = (a, i)$. With this notation, H_S can be written as

$$H_S(v) = -\frac{1}{2}\sum_I \sum_J W_{I,J} v_I v_J + \kappa c \sum_I v_I.$$

Let $v \in (\{0,1\}^l)^c$ not be a fixed point of S. Then there exists K such that $S_K(v) := (S(v))_K \neq v_K$. Then we can define $w \in (\{0,1\}^l)^c$ such that $w_K = S_K(v)$ and $w_I = v_I$ for all $I \neq K$. Then we need to prove that

$$H_S(w) \leq H_S(v).$$

First suppose that $v_K = 0$ and $w_K = S_K(v) = 1$. Then

$$\begin{aligned} & H_S(w) - H_S(v) \\ = \; & -\frac{1}{2}\sum_{I(\neq K)} W_{I,K} v_I - \frac{1}{2}\sum_{J(\neq K)} W_{K,J} v_J \\ & +\kappa c \\ = \; & -\sum_I W_{I,K} v_I + \kappa c \\ \leq \; & 0 \end{aligned}$$

Here we first use the symmetry $W_{I,J} = W_{J,I}$, and we exploit $W_{K,K} = 0$. Then finally, the last inequality is true since $S_K(v) = 1$ if and only if $\sum_I W_{I,K} v_I - \kappa c \geq 0$.

The second case is very similar : if $v_k = 1$ and $w_K = S_K(v) = 0$, then

$$\begin{aligned} & H_S(w) - H_S(v) \\ = \; & \frac{1}{2}\sum_{I(\neq K)} W_{I,K} v_I \\ & +\frac{1}{2}\sum_{J(\neq K)} W_{K,J} v_J - \kappa c \\ = \; & \sum_I W_{I,K} v_I - \kappa c \\ < \; & 0 \end{aligned}$$

Here we use that $W_{K,K} = 0, v_K = 1$ and the last inequality is true since $S_K(v) = 0$ if and only if $\sum_I W_{I,K} v_I - \kappa c < 0$.

We deduce that $H_S(v(t+1)) \leq H_S(v(t))$ for all t and that for all initial state $v(0)$, the dynamics converges to configurations minimizing at least locally the Hamiltonian. Moreover, the dynamics converges finally to a fixed point, that is v such that $S(v) = v$, since flipping from 1 to 0 strictly decreases the Hamiltonian and there are only finitely many possible energies. ■

We will now consider the parallel dynamics. Let $T = (\varphi_{(1,1)}, \ldots, \varphi_{(l,c)})$, and for all $v(0)$, the parallel dynamics $v(t+1) = T(v(t))$, i.e. for all $(a,i), v_{(a,i)}(t+1) = \varphi_{(a,i)}(v(t))$. We have the following result :

Proposition III.2: Let $H = H_T$ be the following function:

$$\begin{aligned} H_T(v) \quad = \quad & -\sum_{(a,i)}\sum_{(b,j)} W_{(a,i),(b,j)} v_{(a,i)}\; y_{(b,j)} \\ & +\kappa c \sum_{(a,i)} (v_{(a,i)} + y_{(a,i)}), \end{aligned}$$

where $y = T(v)$. Again H can be considered as the Hamiltonian that turns T into a Hamiltonian dynamics.

Then for all M, c, l, $v(0)$,

$$H_T(v(t+1)) \leq H_T(v(t)) \qquad \text{for all } t,$$

and the parallel dynamics converges to a fixed point v^*, i.e. $T(v^*) = v^*$, or to a limit cycle $(\tilde{v}^1, \tilde{v}^2)$ of length 2, i.e.

$$T(\tilde{v}^1) = \tilde{v}^2 \quad \text{and } T(\tilde{v}^2) = \tilde{v}^1.$$

Proof: With the previous notations, H_T can be written as

$$\begin{aligned} H_T(v(t)) \quad = \quad & -\sum_I \sum_J W_{I,J} v_I(t)\; v_J(t+1) \\ & +\kappa c \sum_I (v_I(t) + v_I(t+1)) \end{aligned}$$

and

$$\begin{aligned} & H_T(v(t+1)) - H_T(v(t)) \quad = \\ & -\sum_J (v_J(t+2) - v_J(t))(\sum_I W_{J,I} v_I(t+1) - \kappa c), \end{aligned}$$

since W is symmetric. And thus,

$$H_T(v(t+1)) - H_T(v(t)) \leq 0.$$

Indeed,

- if $\sum_I W_{J,I} v_I(t+1) - \kappa c \geq 0$, then $v_J(t+2) = 1$, and $v_J(t+2) - v_J(t) \geq 0$,

- if $\sum_I W_{J,I} v_I(t+1) - \kappa c < 0$, then $v_J(t+2) = 0$, and $v_J(t+2) - v_J(t) \leq 0$.

This proves that H_T is decreasing along each trajectory of the parallel dynamics and this dynamics converges to a subset V such that $H_T(T(v)) - H_T(v) = 0$, for all $v \in V$. Moreover, if we consider the 2-steps dynamics $v(t+2) = T \circ T(v(t))$, the previous calculations show that, there exists a t_0 such that for all $t \geq t_0$, for all J,

- either $v_J(t+2) = v_J(t)$,
- or $v_J(t+2) \neq v_J(t)$ and

$$\sum_I W_{J,I} v_I(t+1) - \kappa c = 0,$$

which implies $v_J(t) = 0, v_J(t+2) = 1$.
Thus, after a finite number of steps, we have $v_I(t+2) = v_I(t)$, for all I, i.e. the dynamics converges to a fixed point or to a 2-cycle.

■

IV. Bounds on the storage capacity

The purpose of the present section is to analyze whether one can find threshold κ such that an amount of $M = \alpha l^2$ patterns can be stored in the above described variant of the GB model. As we will see that this is indeed the case: Also in this sense our model is considerably more effective than the Hopfield model. Recall that in the latter, only for $M \leq N/(2 \log N)$ all patterns are fixed points of the retrieval dynamics (see [15], [5],[4],and [12], [13] for correlated information). In our variant of the GP model there are $N = l \log l$ neurons, while the number of messages to be stored is proportional to l^2. More precisely, we will prove the following theorem.

Theorem IV.1: In the above model with coding matrix W, let $c = \log l$ and $M = \alpha l^2$. For $\kappa \leq 1 - 1/c$ we have
1. If $\alpha < \kappa$, for every fixed μ, every fixed block a and every fixed coordinate k, we have that

$$\mathbb{P}\left(\varphi_{(a,k)}(m^\mu) \neq m^\mu_{(a,k)}\right) \to 0$$

as l and therefore N tends to infinity.
2. If $\alpha < \kappa \exp(-(3+\kappa)/\kappa)$, we have that

$$\mathbb{P}\left(\exists \mu \in \{1, \ldots, M\} : T(m^\mu) \neq m^\mu\right) =$$
$$\mathbb{P}\left(\exists a, \exists i, \exists \mu : \varphi_{(a,i)}(m^\mu) \neq m^\mu_{(a,i)}\right) \to 0$$

as l and therefore N tends to infinity.

Remark IV.2: Again, the comparison to the Hopfield model is interesting. As mentioned above the storage capacity with N neurons there is $N/(2 \log N)$. However, note that the information in our case is extremely sparse. Indeed, only $c = \log l$ out of the $l \log l$ many entries of a message are non-zero. A similar situation for the Hopfield model was considered in the non-rigorous paper by Amari [7]. He asserts that for a Hopfield model with N neurons, where each message has $\mathcal{O}(\log N)$ many 1's and the other neurons are 0, the storage capacity is of order $N^2/(\log N)^2$, which is only slightly worse than our result.

Moreover, we will also be interested in the error correcting abilities of the model. Such errors in a message occur, if some characters are false, or erased. In fact, both types of error are equivalent, if we replace each missing character with a randomly chosen letter. If we not only require the messages to be fixed points of the network dynamics, but also that the network is able to correct a certain percentage of errors (hence works as a truly associative memory), this may lower its capacity. However, the order of the capacity is maintained as can be read off from the following theorem. We will concentrate on one step of the parallel dynamics. To this end we define the discrete ball of radius r centered in m^μ as

$$\begin{aligned} \mathcal{B}(m^\mu, r) &= \{m \in \mathcal{A}^c : d_H(m^\mu, m) \\ &\quad = \text{card}\{j : m^\mu_j \neq m_j\} \leq r\}. \end{aligned}$$

With this notation we have the following result.

Theorem IV.3: In the above model with coding matrix W, let $c = \log l$ and $M = \alpha l^2$. Moreover choose $\gamma \in]0, 1[$, and set $\kappa = \min\{1 - \gamma, 1 - 1/c\}$, and take $\alpha < \kappa \exp(-(1+\kappa)/\kappa))$.

Then for all $\mu = 1, \ldots, M$ and v randomly chosen in $\mathcal{B}(m^\mu, \gamma c - 1)$, we have

$$\mathbb{P}\left(T(v) = m^\mu\right) \to 1$$

as l and therefore N tends to infinity.

In other words for $\alpha > 0$ small enough we can correct any randomly chosen input with at most $\gamma c - 1$ errors.

Remark IV.4: Observe that, other than in the Hopfield model, our model can even repair more than 50 percent mistakes. This is due to the fact that in the Hopfield model the situation between inputs $+1$ and -1 is symmetric, while here a signal 0 is something completely different than a signal 1. Repairing wrong input messages is also a problem in the original GB model, since there the number of 1's is decreasing with the dynamics, i.e. a wrong 0 will always stay a wrong 0.

Let us remark that the Theorem IV.3 is true for random errors and can't be extended to all messages in $\mathcal{B}(m^\mu, \gamma c - 1)$, for $\gamma \in]0,1[$. For example, if we consider the message v, which coincides with the message m^1 for the blocks $a = 1, \dots, c/2$ and with m^2 for the block $c/2+1, \dots, c$, we have necessarily $T(v) \neq m^1$ or $T(v) \neq m^2$. The correction of all possible corrupted messages with a given number of errors is a really different problem: as mentioned in [9] for the Hopfield model, random errors and worst case errors behave in entirely different ways.

For the Hopfield model, using negative association, Bovier proved in [4] that the capacity $M = N/(2\log N)$ is in a certain sense optimal and cannot be surpassed. For our model, Theorem IV.1 states that for $M = \alpha l^2$, with α small enough, the memorized messages are fixed points of the dynamics: From the proof of this theorem in the following section, we see that if $\kappa \leq 1 - 1/c$ and $\alpha \leq \kappa \exp(-(1+\kappa)/\kappa)$, then for all $\mu = 1, \dots, M$,

$$\lim_{l\to\infty} \mathbb{P}(T(m^\mu) = m^\mu) = 1.$$

In particular, if we choose $\kappa = 1 - 1/c$, we obtain the maximal value

$$\alpha^* = (1-1/c)\exp(-1-\frac{c}{c-1}) \simeq e^{-2} \simeq 0.135.$$

A natural question is whether the bound we get for α is optimal. A first step in this direction is the following result where we obtain an upper bound for the capacity, in the sense of the stability of the memorized messages.

Theorem IV.5: Let $\kappa \leq 1 - 1/c$. Then, for each $\alpha > -\log(1 - e^{-1}) \simeq 0.45$ and each $\mu \in \{1, \dots, M\}$, the message m^μ is not stable with probability converging to 1 when l tends to infinity, i.e.

$$\lim_{l\to\infty} \mathbb{P}(T(m^\mu) \neq m^\mu) = 1.$$

Remark IV.6: All the results concerning the stability of stored messages are the same if we consider the sequential dynamics S instead of the parallel dynamics T.

V. Proof of Theorems IV.1 and IV.3

In this section Theorems IV.1 and IV.3 are proved. The proofs are rather similar and employ exponential inequalities. We will treat the proof of Theorem IV.1 in greater details and will be slightly more sketchy about the proof of Theorem IV.3.

Proof: [Proof of Theorem IV.1] For a learned message $m^\mu \in \mathcal{M}$ we have to check, whether m^μ is a fixed point of the dynamics φ. Without loss of generality $\mu = 1$ and the first pattern has a one entry at the first position of every block and 0's otherwise. That is to say $m^1 = (e_1, \dots, e_1)$.

Now there are two cases to consider, depending on whether we want to check that a 1 remains a 1 under the dynamics or a 0 remains a 0. Let us start with the first case.

For $i = 1$, we have for $1 \leq a \leq c$

$$\sum_{b=1}^{c}\sum_{j=1}^{l} W_{(a,1),(b,j)} m^1_{(b,j)} = \sum_{b=1}^{c} W_{(a,1),(b,1)} \geq c-1,$$

so if we choose $\kappa \leq 1 - 1/c$, we immediately see that

$$\mathbb{P}(\varphi_{(a,i)}(v) = 1) = 1.$$

So all entries with a 1 are trivially fixed points of our dynamics.

For the case that the entry of m^1 is zero, we may without loss of generality consider the

case $a = 1$ and $i = 2$. Then

$$\begin{aligned}
&\mathbb{P}\left(\varphi_{(1,2)}(m^1) \neq m^1_{(1,2)}\right)\\
=\;& \mathbb{P}\left(\sum_{b=2}^{c}\sum_{j=1}^{l} W_{(1,2),(b,j)} m^1_{(b,j)} \geq \kappa c\right)\\
=\;& \mathbb{P}\left(\sum_{b=2}^{c}\sum_{j=1}^{l}\sum_{\mu=2}^{M} \zeta^\mu_{(1,2)}\zeta^\mu_{(b,j)} m^1_{(b,j)} \geq \kappa c\right)\\
=\;& \mathbb{P}\left(\sum_{b=2}^{c}\sum_{\mu=2}^{M} \zeta^\mu_{(1,2)}\zeta^\mu_{(b,1)} \geq \kappa c\right)\\
\leq\;& e^{-t\kappa c}\mathbb{E}\left[e^{t\sum_{b=2}^{c}\sum_{\mu=2}^{M}\zeta^\mu_{(1,2)}\zeta^\mu_{(b,1)}}\right]\\
=\;& e^{-t\kappa c}\mathbb{E}\left[e^{t\sum_{b=2}^{c}\zeta^2_{(1,2)}\zeta^2_{(b,1)}}\right]^{M-1} \qquad (2)
\end{aligned}$$

for some $t > 0$. Here we just apply the definition of the dynamics, the definition of W, the fact that we know, for which coordinates m^1 has a non-zero entry, and finally an exponential Chebyshev-inequality and the independence of the messages.

We compute $\mathbb{E}\left[e^{t\sum_{b=2}^{c}\zeta^2_{(1,2)}\zeta^2_{(b,1)}}\right]$ and obtain the following bound:

$$\begin{aligned}
&\mathbb{E}\left[e^{t\sum_{b=2}^{c}\zeta^2_{(1,2)}\zeta^2_{(b,1)}}\right]\\
=\;& \left(1-\frac{1}{l}\right)\cdot 1 + \frac{1}{l}\left(\mathbb{E}[e^{t\sum_{b=2}^{c}\zeta^2_{(b,1)}}]\right)\\
\leq\;& \left(1-\frac{1}{l}\right) + \frac{1}{l}\left(\frac{e^t-1}{l}+1\right)^c\\
\leq\;& \left(1-\frac{1}{l}\right) + \frac{1}{l}e^{c\cdot\frac{e^t-1}{l}},
\end{aligned}$$

where we used the standard estimate $1+x \leq e^x$ for all x. Plugging this into (2), we arrive at

$$\begin{aligned}
&\mathbb{P}\left(\varphi_{(1,2)}(m^1) \neq m^1_{(1,2)}\right) \qquad (3)\\
\leq\;& e^{-t\kappa c}\left[\left(1-\frac{1}{l}\right) + \frac{1}{l}e^{c\cdot\frac{e^t-1}{l}}\right]^M\\
\leq\;& e^{-t\kappa c}e^{\frac{M}{l}\left(e^{c\cdot\frac{e^t-1}{l}}-1\right)}.
\end{aligned}$$

Now, for fixed t and $l \to \infty$, we have by expanding the first exponential in the exponent

$$\begin{aligned}
&\mathbb{P}\left(\varphi_{(1,2)}(m^1) \neq m^1_{(1,2)}\right)\\
\leq\;& e^{-t\kappa c}\exp\left(\frac{M}{l}\left(e^{c\cdot\frac{e^t-1}{l}}-1\right)\right)\\
=\;& e^{-t\kappa c}e^{\frac{M}{l}\left(c\frac{e^t-1}{l}+\frac{1}{2}\left(c\frac{e^t-1}{l}\right)^2+O\left(\left(c\frac{e^t-1}{l}\right)^3\right)\right)}\\
=\;& e^{-t\kappa c}e^{\frac{Mc}{l^2}(e^t-1)+\frac{Mc^2}{2l^3}(e^t-1)^2+\frac{M}{l}O\left(\left(\frac{c}{l}(e^t-1)\right)^3\right)}\\
\approx\;& \exp\left(c\left(-t\kappa+\alpha(e^t-1)\right)\right),
\end{aligned}$$

where for two sequences (a_l) and (b_l) we write $a_l \approx b_l$, if the fraction a_l/b_l tends to 1 as l tends to infinity.

The right hand side above takes its minimum at $t = \log(\kappa/\alpha)$. Plugging this value in yields

$$\mathbb{P}\left(\varphi_{(1,2)}(m^1) \neq m^1_{(1,2)}\right) \leq l^{\kappa-\alpha-\kappa\log(\kappa/\alpha)}.$$

If now $\alpha < \kappa$ the exponent is negative. This proofs the first assertion of the Theorem.

Moreover,

$$\begin{aligned}
&\mathbb{P}\left(\exists 1 \leq a \leq c, 1 \leq i \leq l, 1 \leq \mu \leq M : \varphi_{(a,i)}(m^\mu) \neq m^\mu_{(a,i)}\right)\\
&\qquad \leq Mlc\, e^{c\left(-t\kappa+\alpha(e^t-1)\right)}
\end{aligned}$$

Due to our choice of $M = \alpha l^2$ and $c = \log l$ this converges to 0, if

$$-t\kappa + \alpha(e^t-1) < -3.$$

If we replace t by $\log(\kappa/\alpha)$ we obtain

$$-\kappa\log\left(\frac{\kappa}{\alpha}\right) + \kappa - \alpha < -3.$$

We deduce that the previous inequality is true, if

$$-\kappa\log\left(\frac{\kappa}{\alpha}\right) < -3-\kappa,$$

that is $\alpha < \kappa\exp(-(3+\kappa)/\kappa)$.

■

We will now continue with proof of Theorem IV.3.

Proof: [Proof of Theorem IV.3] Choose $0 < \gamma < 1$ and as indicated in the theorem, we set $\kappa = 1-\gamma$. Suppose we want to correct a message with r errors, where $r \leq \gamma c - 1$. Again without loss of generality we concentrate on the case $\mu = 1$ and $m^1 = (e_1, \ldots, e_1)$. Let $v = (e_2, \ldots, e_2, e_1, \ldots, e_1)$ be the vector consisting

of e_2 in the first r places and e_1 in the last $c-r$ places, i.e. v is m^1 with r errors.

Now there are now four different cases to differentiate:

- The false 1's in the clusters with wrong entries have to be turned into a 0,
- The correct 0's in all clusters have to remain 0s,
- The wrong 0's have to be turned into 1's,
- and the correct 1's have to remain 1's.

The last two cases indeed are similar and with our choice of κ, they have probability one. Indeed, for example take $a=1$, $b=1$. Using $v_{(b,j)}=1$ only if $b\leq r, j=2$, or $b\geq r+1, j=1$, we obtain

$$\begin{aligned}
&\mathbb{P}\left(\sum_{b=1}^{c}\sum_{j=1}^{l} W_{(1,1),(b,j)}v_{(b,j)} < \kappa c\right)\\
=\;&\mathbb{P}\left(\sum_{b=1}^{r}\sum_{j=1}^{l} W_{(1,1),(b,j)}v_{(b,j)} \right.\\
&\qquad\left. < \kappa c - \sum_{b=r+1}^{c} W_{(1,1),(b,1)}\right)\\
\leq\;&\mathbb{P}\left(\sum_{b=1}^{r}\sum_{j=1}^{l} W_{(1,1),(b,j)}v_{(b,j)} < \kappa c - c + r\right)\\
=\;&0,
\end{aligned}$$

since $\sum_{b=r+1}^{c} W_{(1,1),(b,1)} \geq c-r$ and $\kappa c - c + r \leq -\gamma c + r \leq 0$.

This proves that a wrong 0 becomes a 1 with probability 1. Let us now consider the case of the correct 1's. For example, for $c=1$ and $b=1$, we obtain

$$\begin{aligned}
&\mathbb{P}\left(\sum_{b=1}^{c}\sum_{j=1}^{l} W_{(c,1),(b,j)}v_{(b,j)} < \kappa c\right)\\
=\;&\mathbb{P}\left(\sum_{b=1}^{r}\sum_{j=1}^{l} W_{(c,1),(b,j)}v_{(b,j)} \right.\\
&\qquad\left. < \kappa c - \sum_{b=r+1}^{c} W_{(c,1),(b,1)}\right)\\
\leq\;&\mathbb{P}\left(\sum_{b=1}^{r}\sum_{j=1}^{l} W_{(1,1),(b,j)}v_{(b,j)} \right.\\
&\qquad\left. < \kappa c - c + r + 1\right)\\
=\;&0,
\end{aligned}$$

since $\sum_{b=r+1}^{c} W_{(c,1),(b,1)} \geq c-r-1$ and $\kappa c - c + r - 1 \leq -\gamma c + r - 1 \leq 0$.

The first two cases are also identical and can be treated analogously to our last calculations: For the first case and $a=1$, using $\zeta^1_{(1,2)}=0$, we get exemplary

$$\begin{aligned}
&\mathbb{P}\left(\sum_{b=1}^{c}\sum_{j=1}^{l} W_{(1,2),(b,j)}v_{(b,j)} \geq \kappa c\right)\\
=\;&\mathbb{P}\left(\sum_{b=2}^{c}\sum_{j=1}^{l}\sum_{\mu=2}^{M} \zeta^{\mu}_{(1,2)}\zeta^{\mu}_{(b,j)}v_{(b,j)} \geq \kappa c\right)\\
=\;&\mathbb{P}\left(\sum_{b=2}^{r}\sum_{\mu=2}^{M} \zeta^{\mu}_{(1,2)}\zeta^{\mu}_{(b,2)} + \sum_{b=r+1}^{c}\sum_{\mu=2}^{M} \zeta^{\mu}_{(1,2)}\zeta^{\mu}_{(b,1)} \geq \kappa c\right)\\
\leq\;&e^{-t\kappa c}\mathbb{E}\left[e^{t\sum_{b=2}^{c}\sum_{\mu=2}^{M}\zeta^{\mu}_{(1,2)}\zeta^{\mu}_{(b,1)}}\right]\\
\leq\;&e^{c\left(-t\kappa+\alpha(e^t-1)\right)}
\end{aligned}$$

This function is minimal in $t=\log(\kappa/\alpha)$.

So, if we want to correct a randomly chosen message $v\in\mathcal{B}(m^\mu,\gamma c-1)$, we need to bound the following probability

$$\begin{aligned}
\mathbb{P}\left(T(v)\neq m^\mu\right) &\leq cle^{c\left(-t\kappa+\alpha(e^t-1)\right)}\\
&\leq e^{c\left(-t\kappa+\alpha(e^t-1)+1+o(1)\right)},
\end{aligned}$$

We want that the right hand side converges to 0, which is the case if

$$-t\kappa+\alpha(e^t-1)<-1.$$

After filling in the minimizing t, i.e. $t = \log(\kappa/\alpha)$, this becomes

$$-\log\left(\frac{\kappa}{\alpha}\right)\kappa + \kappa - \alpha < -1,$$

which is true if

$$\alpha < \kappa \exp(-\frac{1+\kappa}{\kappa}),$$

with $\kappa = 1 - \gamma$. This finishes the proof of Theorem IV.3. ∎

VI. Proof of Theorem IV.5

In this section, we prove the lower bound on the storage capacity. i.e. Theorem IV.5.

Proof: [Proof of Theorem IV.5] Again, without loss of generality, we take $\mu = 1$ and $m^\mu = (e_1, \dots, e_1)$ Observe that the message m^1 cannot be a fixed point of the dynamics, if there is at least one node $v_{(a,i)}$, $i \neq 1$, with $\varphi_{(a,i)}(m^1) = 1$, which holds if

$$\sum_{b=1}^{c} W_{(a,i),(b,1)} v_{(b,1)} \geq c - 1,$$

and

$$\mathbb{P}(T(m^1) \neq m^1) \quad = \\ \mathbb{P}(\exists a \in \{1,\dots,c\}, \exists i \in \{2,\dots,l\} : \varphi_{(a,i)}(m^1) = 1).$$

This is in particular the case if $v_{(a,i)}$ is connected to every neuron of m^1,

$$\begin{aligned} & \mathbb{P}(T(m^1) \neq m^1) \\ \geq\ & \mathbb{P}(\exists a \in \{1,\dots,c\}, \exists i \in \{2,\dots,l\}, \\ & \quad \forall b \in \{1,\dots,c\} \setminus \{a\} : W_{(a,i),(b,1)} \geq 1). \end{aligned}$$

This probability can be bounded by

$$\begin{aligned} & \mathbb{P}(T(m^1) \neq m^1) \\ \geq\ & \mathbb{P}(\exists a \in \{1,\dots,c\}, \exists i \in \{2,\dots,l\}, \\ & \quad \forall b \in \{1,\dots,c\} \setminus \{a\} : W_{(a,i),(b,1)} \geq 1) \\ =\ & 1 - \mathbb{P}(\forall a \in \{1,\dots,c\}, \forall i \in \{2,\dots,l\} \\ & \quad \exists b \in \{1,\dots,c\} \setminus \{a\} : W_{(a,i),(b,1)} < 1) \\ \geq\ & 1 - \mathbb{P}(\forall i \in \{2,\dots,l\} \exists b \in \{2,\dots,c\} : \\ & \quad W_{(1,i),(b,1)} < 1). \end{aligned}$$

The neurons $(1,i)$ and $(1,j)$, $i \neq j$, cannot be part of the same message. Hence, for $i \neq j$, we have

$$\begin{aligned} & \mathbb{P}(\exists b_1, b_2 \in \{2,\dots,c\} : \\ & \quad W_{(1,i),(b_1,1)} < 1, W_{(1,j),(b_2,1)} < 1) \\ \leq\ & \mathbb{P}(\exists b_1 \in \{2,\dots,c\} : W_{(1,i),(b_1,1)} < 1) \times \\ & \quad \mathbb{P}(\exists b_2 \in \{2,\dots,c\} : W_{(1,j),(b_2,1)} < 1). \end{aligned}$$

Instead of an exact formal proof we give a more intuitive justification: If $(1,i)$ is not connected to every neuron m^1_a of the other clusters a, $2, \leq a \leq c$, it will be more likely (but definitely not less likely, this is all we need) that $(1,j)$ is connected to every neuron m^1_a of the other clusters, because $(1,i)$ and $(1,j)$ cannot be part of the same message. This means on the one hand that the fact that $(1,i)$ is not connected to every other cluster does not give any information about the neurons m^1_a being activated in the same messages than $(1,j)$. On the other hand, $(1,i)$ and $(1,j)$ share the first cluster's $M-1$ places in the $M-1$ remaining messages, so they compete for the places, and the fact that $(1,i)$ is not connected to each of the other clusters indicates less messages containing $(1,i)$ and potentially more that can contain $(1,j)$. This argument can in fact be turned into a rigorous proof.
Altogether we have

$$\begin{aligned} & \mathbb{P}(\forall i \in \{2,\dots,l\} \exists b \in \{2,\dots,c\} : W_{(1,i),(b,1)} < 1) \\ \leq\ & \prod_{i=2}^{l} \mathbb{P}(\exists b \in \{2,\dots,c\} : W_{(1,i),(b,1)} < 1) \\ =\ & \mathbb{P}(\exists b \in \{2,\dots,c\} : W_{(1,2),(b,1)} < 1)^l. \end{aligned}$$

Thus, we have to estimate

$$\mathbb{P}(\exists b \in \{2,\dots,c\} : W_{(1,2),(b,1)} < 1).$$

Lemma VI.1: Let $m^1 = (e_1, \dots, e_1)$, $a \in \{1,\dots,c\}$ and $i \in \{2,\dots,l\}$.
We define Y by the number of neurons belonging to the message m^1 which are connected to (a,i). Then for the distribution of Y, it holds that

$$\mathbb{P}(Y = i) = \binom{c-1}{i} \left(1 - e^{-\alpha}\right)^i e^{-\alpha(c-1-i)}(1+o(1))$$

for $i \in \{0,\dots,c-1\}$, and 0, otherwise. In

particular,

$$\begin{aligned}&\mathbb{P}(\forall b \in \{1,\dots,c\} \setminus \{a\} : W_{(a,i),(b,1)} = 1)\\ =\;& 1 - \left(1-e^{-\alpha}\right)^{c-1}(1+o(1))\end{aligned}$$

as $l \to \infty$.

With Lemma VI.1 in hands we can conclude the Proof of Theorem IV.5 in the following way:

$$\begin{aligned}&\mathbb{P}(\exists b \in \{2,\dots,c\} : W_{(1,2),(b,1)} < 1)\\ =\;& 1 - \mathbb{P}(\forall b \in \{2,\dots,c\} : W_{(1,2),(b,1)} = 1)\\ =\;& 1 - \left(1-e^{-\alpha}\right)^{c-1}(1+o(1)).\end{aligned}$$

Finally,

$$\begin{aligned}&\mathbb{P}(T(m^1) \neq m^1)\\ &\geq 1 - \left(1-(1-e^{-\alpha})^{c-1}(1+o(1))\right)^{l} \overset{l\to\infty}{\to} 1\end{aligned}$$

for $\alpha > -\log(1-e^{-1}) \approx 0.45$. ■

It remains to prove Lemma VI.1.

Proof: [Proof of Lemma VI.1] We choose without loss of generality $a = 1$ and $i = 2$. Let Y be the number of neurons belonging to the message m^1 which are connected to $(1,2)$. To determine the distribution of Y, we split the event $\{Y = i\}$ into the disjoint events

$$\{Y = i, Z_{(1,2)} = n\},$$

where $Z_{(1,2)}$ is the number of messages containing $(1,2)$.

Then,

$$\begin{aligned}&\mathbb{P}(Y = i | Z_{(1,2)} = n)\\ =\;& \left(1-\left(1-\frac{1}{l}\right)^{n}\right)^{i}\left(1-\frac{1}{l}\right)^{n(c-1-i)}.\end{aligned}$$

We arrive at

$$\begin{aligned}&\mathbb{P}(Y = i)\\ =\;& \binom{c-1}{i}\sum_{n=0}^{M-1}\binom{M-1}{n}\frac{1}{l^n}\left(1-\frac{1}{l}\right)^{M-1-n}\\ &\times\left(1-\left(1-\frac{1}{l}\right)^{n}\right)^{i}\left(1-\frac{1}{l}\right)^{n(c-1-i)}\\ =\;& \binom{c-1}{i}\sum_{n=0}^{M-1}\binom{M-1}{n}\frac{1}{l^n}\\ &\times\left(1-\frac{1}{l}\right)^{M-1-n+n(c-1-i)}\\ &\times\sum_{k=0}^{i}\binom{i}{k}(-1)^k\left(1-\frac{1}{l}\right)^{nk}\\ =\;& \binom{c-1}{i}\sum_{k=0}^{i}\binom{i}{k}(-1)^k\sum_{n=0}^{M-1}\\ &\times\binom{M-1}{n}\frac{1}{l^n}\left(1-\frac{1}{l}\right)^{M-1-n+nk+n(c-1-i)}\\ =\;& \binom{c-1}{i}\sum_{k=0}^{i}\binom{i}{k}(-1)^k\\ &\times\left(\frac{1}{l}\left(1-\frac{1}{l}\right)^{k+c-1-i}+1-\frac{1}{l}\right)^{M-1}\\ =\;& \binom{c-1}{i}\sum_{k=0}^{i}\binom{i}{k}(-1)^k\\ &\times\left(\frac{1}{l}-\frac{k+c-1-i}{l^2}+O(l^{-3})+1-\frac{1}{l}\right)^{M-1}\\ =\;& \binom{c-1}{i}\sum_{k=0}^{i}\binom{i}{k}(-1)^k\\ &\times\left(1-\frac{k+c-1-i}{l^2}+O(l^{-3})\right)^{M-1}\\ =\;& \binom{c-1}{i}\sum_{k=0}^{i}\binom{i}{k}(-1)^k\\ &\times e^{-\alpha(k+c-1-i)}(1+o(1))\\ =\;& \binom{c-1}{i}\left(1-e^{-\alpha}\right)^{i}e^{-\alpha(c-1-i)}(1+o(1)).\end{aligned}$$

In particular

$$\mathbb{P}(Y = c-1) = \left(1-e^{-\alpha}\right)^{c-1}(1+o(1)).$$

■

References

[1] B. K. Aliabadi, C. Berrou, V. Gripon, and X. Jiang. Learning sparse messages in networks of neural cliques. *IEEE Transactions on Neural Networks and Learning Systems*, August 2012. To appear.

[2] B. K. Aliabadi, C. Berrou, V. Gripon, and X. Jiang. Storing sparse messages in networks of neural cliques. *IEEE Transactions on Neural Networks and Learning Systems*, 25:980–989, 2014.

[3] R. B. Ash. *Information theory*. Dover Publications, Inc., New York, 1990. Corrected reprint of the 1965 original.

[4] A. Bovier. Sharp upper bounds on perfect retrieval in the Hopfield model. *J. Appl. Probab.*, 36(3):941–950, 1999.

[5] D. Burshtein. Nondirect convergence radius and number of iterations of the Hopfield associative memory. *IEEE Trans. Inform. Theory*, 40(3):838–847, 1994.

[6] V. Gripon and C. Berrou. Sparse neural networks with large learning diversity. *IEEE Transactions on Neural Networks*, 22(7):1087–1096, July 2011.

[7] S. ichi Amari. Characteristics of sparsely encoded associative memory. *Neural Networks*, 2(6):451–457, 1989.

[8] A. Jagota, G. Narasimhan, and K. W. Regan. Information capacity of binary weights associative memories. *Neurocomputing*, 19:35–58, 1996.

[9] J. Komlós and R. Paturi. Convergence results in an associative memory model. *Neural Networks*, 1(3):239–250, 1988.

[10] D. Loukianova. Lower bounds on the restitution error in the Hopfield model. *Probab. Theory Related Fields*, 107(2):161–176, 1997.

[11] M. Löwe. On the storage capacity of Hopfield models with correlated patterns. *Ann. Appl. Probab.*, 8(4):1216–1250, 1998.

[12] M. Löwe. On the storage capacity of Hopfield models with correlated patterns. *Ann. Appl. Probab.*, 8(4):1216–1250, 1998.

[13] M. Löwe and F. Vermet. The storage capacity of the Hopfield model and moderate deviations. *Statist. Probab. Lett.*, 75(4):237–248, 2005.

[14] M. Löwe and F. Vermet. The Hopfield model on a sparse Erdős-Renyi graph. *J. Stat. Phys.*, 143(1):205–214, 2011.

[15] R. J. McEliece, E. C. Posner, E. R. Rodemich, and S. S. Venkatesh. The capacity of the Hopfield associative memory. *IEEE Trans. Inform. Theory*, 33(4):461–482, 1987.

[16] C. M. Newman. Memory capacity in neural network models: Rigorous lower bounds. *Neural Networks*, 1(3):223–238, 1988.

[17] M. Talagrand. Rigorous results for the Hopfield model with many patterns. *Probab. Theory Related Fields*, 110(2):177–276, 1998.